Omni-MVA
Network Switching System

PRODUCT REQUIREMENTS DOCUMENT

REVISION HISTORY

	Date	Revision Description	Changed by:	Approved by
	4/24/96	Initial Release for comment	G. Stone P. Terry C. Haywood	

Table of Contents

1 OVERVIEW 1

2 APPLICATIONS ..... Error! Bookmark not defined.

3 XYLAN POSITIONING 1

4 FUNCTIONAL OVERVIEW ..... Error! Bookmark not defined.

4.1 PHYSICAL PACKAGING 1
4.2 ARCHITECTURE 1

5 COST TARGETS ..... Error! Bookmark not defined.

6 AVAILABILITY 5

1 Overview

This document describes the functional requirements of the Omni-MVA switch. The Omni-MVA switch is a follow-on to the OmniSwitch product with greatly increased bandwidth, port density and switching performance over the existing OmniSwitch and stackable products.

Omni-MVA utilizes a 10.24 gigabit non-blocking frame switching fabric, with a peak performance of 6.6 million frames per second. Individual switching modules are allotted 1.28 gigabits of bandwidth.

Omni-MVA uses the Omni 5X and 9X backplanes, so that existing chassis can be upgraded to Omni-MVA. Existing OmniSwitch modules can be used in conjunction with Omni-MVA. As part of this development, a new chassis is planned that increases the width of slots to 1.4 inches, allowing higher port density on switching modules. These new modules will not be useable in the older chassis.

2 Xylan Positioning

The Omni-MVA switch is a direct response to the accelerating acceptance of 100 megabit Ethernet, as well as being a response to certain competitive products now in development. At introduction, Omni-MVA will be the highest performance LAN switch available on the market. It will also have the lowest price-per-port of any full-function LAN switch.

The Omni-MVA is designed to provide non-blocking, wire-rate switching for up to 64 100BaseTx LANs, or up to 256 10BaseT ports.

3 Architecture

The Omni-MVA architecture consists of a high performance, non-blocking, frame switching fabric interconnecting frame switching modules. Each switching module has 1.28 gigabits of bandwidth into the fabric. Total fabric bandwidth is 10.24 gigabits per second.

A new high-performance CAM enables a frame switching performance of 6.6 million frames per second.

All Omni-MVA modules are derived from a common set of ASIC building blocks. This allows individual modules to be tailored to specific cost/performance targets while nearly eliminating the need for additional ASIC development

3.1 System structure

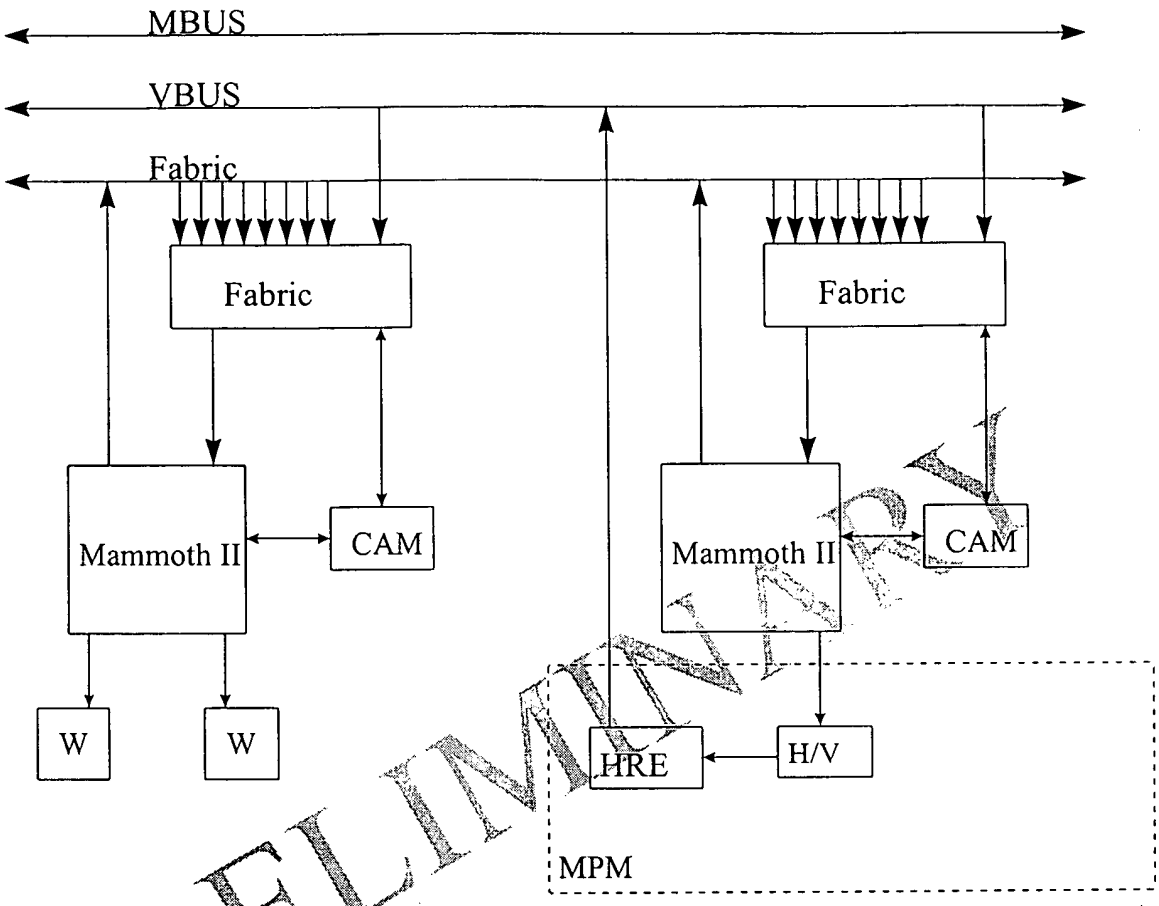

3.2 Switching Fabric

The Omni-MVA switching fabric operates as eight parallel VBUS links, operating at 1.28 gigabits each. Each switching module is able to transmit on a single link, while all modules are able to receive simultaneously on all links.

The frame format used on the link is identical to VBUS. When a switching modules transmits on a link, all modules perform a CAM lookup and execute the claiming logic in the manner of VBUS.

3.3 ASIC set

Each of the different types of modules used in Omni-MVA are derived from a basic set of hardware building blocks each of which is an individual ASIC. This approach results in a high degree of flexibility when individual modules are designed.

The set consists of FABRIC-II, Mammoth-II, CAM, Whistler, and Whistler 1k. Of these ASICs, all are new for the Omni-MVA except for the Whistler, which was developed for the Mammoth-based stackables.

3.3.1 FABRIC-II

The FABRIC-II ASIC contains that drivers that form the communication links that make up the switching fabric, as well as the logic that performs the VBUS-style claiming functions.

3.3.1.1 IOP

The IOP provides buffer management for native ATM interfaces. It has no SAR capability. Each IOP is capable of interfacing

3.3.1.2 MIOP

3.3.1.3 CAM

3.3.1.4 Whistler

3.3.1.5 Whistler 1000

3.3.2 ASIC set

Frame switching modules incorporate a SAR function to interface to the backplane. Server modules will be provided to provide higher level functions such as routing, and full RMON support. The diagram below illustrates the architectural concept.

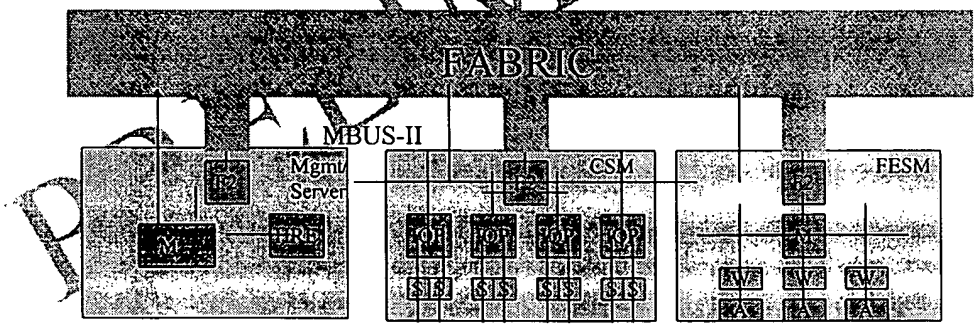

Omni97 Architecture

The diagram above shows the interconnection of the three types of modules. The Fabric speed is (TBD probably about 15 Gbps) and all data between modules crosses the fabric. The MBUS-II performs a similar role to the OmniSwitch MBUS but is cost optimized and provides for multiple bus masters.

The key new element to be developed for the Omni-97 is a new Fabric ASIC (F2) which provides an interface into the Fabric; supports embedded SAR functionality for AAL5; and a frame bus interface to the Mammoth; and cell bus interfaces into the IOPs. This will be a large ASIC and will likely gate the entire project.

The Mgmt/Server module is a module capable of performing as a management processor, route server (layer three switching), ATM call processing server, or RMON server. The multimaster

capability of the MBUS-II will allow for multiple management/server entities to take on independent roles of managing certain functions.

The FESM module is a 24 port 10/100 Ethernet switching module. It is intended that this module will use the Mammoth, Whistler, and Aspen asics. This module will perform local switching between ports and will send frames to the fabric only when needed for inter-module traffic. Several variants of this module will be built, including an HSM-like unit for support of FDDI, Token Ring.

The CSM module is a 8 port OC3 ATM switching module. It is intended that this module will use the IOP asic. A 16 port OC3 UTP, 2 port OC12, and 24 port ATM25 CSM modules are also planned.

A WAN interface module is also planned which will incorporate CBR, Frame and Cell switching functions on a single module.

Since all frame movement between LAN modules through the Fabric will be Cell based, X-LANE, LANE, or MPOA encapsulations need to be supported within the switch as well as outside the switch.

3.4 Features

4 Performance Summary

The target fabric performance is 1.28 gigabits per slot for a total of 10.24 gigabits second total throughput. Each module has access to a total of four 1 gigabit fabric links for a per-slot aggregate of 4 gigabits.

5 Cost Targets

The following table indicates cost targets for the products described in this document. Price targets reflect an 80% gross margin at list and 60% gross margin at 50% discount.

Model	Cost Target*	Price Target (US List)
5 Slot Chassis with Mgmt + 1 PS	1,600	$8,000
13-Slot Chassis with Mgmt + 1 PS	$2,100	$10,500
Route Server	$1,000	$5,000
RMON Server	$1,000	$5,000
Call Processing Server	$1,000	$5,000
ESM - 24	$960	$4,800
FSM - 2	$960	$4,800
CSM25 - 24	$960	4,800
CSMOC3 - 8F	$1920	$9,600
CSMOC12 - 2	$1920	$9,600
CSM25 - 24	$960	4,800

Table 5—Cost and price targets
* Includes 25% manufacturing burden

6 Availability

The modules described in this document are targetted for the following releases:

1.	Omni97-13 Chassis	Release 4.0 (GA Q2 97)
2.	Omni97-5 Chassis	Release 4.1 (GA Q3 97)
3.	10/100 ESM-Tx-24	Release 4.0 (GA Q2 97)
4.	10/100 ESM-Fx-12	Release 4.0 (GA Q2 97)
5.	OC3-UTP-16	Release 4.0 (GA Q2 97)
6.	OC3-UTP-12	Release 4.0 (GA Q2 97)
7.	OC12-SM-4	Release 4.0 (GA Q2 97)
8.	FDDI-2 (HSM/4)	Release 4.1 (GA Q3 97)
9.	Token Ring-12 (HSM/4)	Release 4.1 (GA Q3 97)
10.	Frame Relay-8 (HSM/4)	Release 4.1 (GA Q3 97)
11.	ISDN BRI-4(HSM/4)	Release 4.1 (GA Q3 97)

6.1 Packaging

The Omni-MVA will be implemented with two chassis sizes - a 9 Slot and a 5 slot. All I/O and serviceability will be provided through the front. The packaging concept is illustrated in the drawing below:

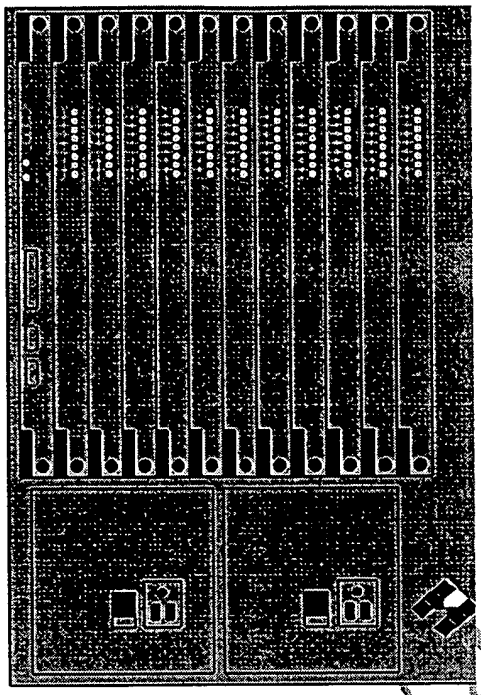

Omni-MVA - 9 Slot Chassis

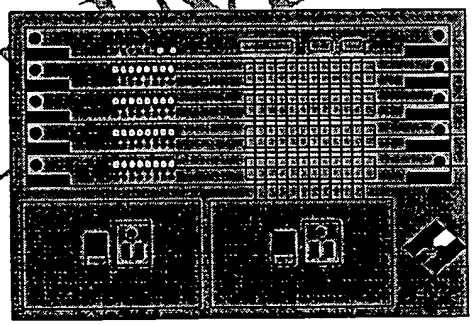

Omni-MVA - 5 Slot Chassis

The Omni-MVA chassis' are required to operate in the same environment as the OmniSwitch - Phone closets and data centers. The Omni-MVA chassis requirements are:

- to be rack mountable in a standard 19" equipment rack
- be wall mountable via optional mounting brackets
- provide service for all cabling and maintenance through the front
- to operate over a temperature range of 15 - 45 degrees C
- provide temperature monitoring to network management
- be made of light weight aluminum

- provide adequate front panel space to support 32 RJ45 connectors

Port count	9 Slot: up to 288 Ethernet or ATM ports 9 Slot: up to 96 Ethernet or ATM ports
Management Interface	One 10 BaseT interface and one, RS-232C, 9-pin "D" connector, configured per IBM AT serial port DTE and one as DCE.
Standards Compliance	IEEE 802.3; CSA 950; CISPR 22, Class A; CISPR 22, Class B (Fiber and STP only); EN50081-1 / EN55022, 1987; EN50082-1, 1992; EN50091-1 / EN55022, 1987; EN60950; FCC Part 15, Subpart B (Class A); IEC 801-2, 1991; IEC 801-3, 1984; IEC 801-4, 1988; UL 1950; VCCI V-3/93.01; preparing for ISO 9000
Power source	90 to 265 VAC, 47 to 63 Hz -48VDC
Ambient operating temperature range	0 to 45 degrees C
Physical dimensions	12 Slot: 17.5" x 21" x 14" 5 Slot: 17.5" x 10.5" x 14"

Omni-MVA Specifications